

IN THE CLAIMS

1. (Currently Amended) A method of mining attribute associations in a relational data set, comprising the steps of:

obtaining multiple items from the relational data set; and

discovering attribute associations using: (i) multi-attribute mining templates formed from at least a portion of the multiple items; and (ii) one or more mining preferences specified by a user;

wherein the multi-attribute mining templates are related by an anti-monotonicity property such that the property holds when mining top-down from k -itemsets to $(k + 1)$ -itemsets and when mining items defined by a set of k attributes to items defined by $k + 1$ attributes.

2. (Canceled).

3. (Original) The method of claim 1, wherein the one or more mining preferences specified by the user comprise specification of at least one of: (i) one or more desired multi-attribute mining templates; (ii) one or more irrelevant multi-attribute mining templates; and (iii) one or more rules concerning values of attributes in the multi-attribute mining templates.

4. (Original) The method of claim 1, wherein the attribute association discovering step further comprises generating candidate patterns at a template level.

5. (Original) The method of claim 4, wherein the candidate pattern generating step further comprises deriving candidate patterns of multi-attribute mining templates by merge-joining patterns of nodes of at least a portion of the templates without pre-sorting.

6. (Original) The method of claim 4, wherein the candidate pattern generating step further comprises maintaining one or more occurrence buffers to count occurrences of patterns.

7. (Original) The method of claim 4, wherein the attribute association discovering step further comprises pruning candidate patterns at a template level.

8. (Currently Amended) Apparatus for mining attribute associations in a relational data set, comprising: a memory; and

at least one processor coupled to the memory and operative to: (i) obtain multiple items from the relational data set; and (ii) discover attribute associations using: (i) multi-attribute mining templates formed from at least a portion of the multiple items; and (ii) one or more mining preferences specified by a user, wherein the multi-attribute mining templates are related by an anti-monotonicity property such that the property holds when mining top-down from k -itemsets to $(k + 1)$ -itemsets and when mining items defined by a set of k attributes to items defined by $k + 1$ attributes.

9. (Canceled).

10. (Original) The apparatus of claim 8, wherein the one or more mining preferences specified by the user comprise specification of at least one of: (i) one or more desired multi-attribute mining templates; (ii) one or more irrelevant multi-attribute mining templates; and (iii) one or more rules concerning values of attributes in the multi-attribute mining templates.

11. (Original) The apparatus of claim 8, wherein the attribute association discovering operation further comprises generating candidate patterns at a template level.

12. (Original) The apparatus of claim 11, wherein the candidate pattern generating operation further comprises deriving candidate patterns of multi-attribute mining templates by merge-joining patterns of nodes of at least a portion of the templates without pre-sorting.

13. (Original) The apparatus of claim 11, wherein the candidate pattern generating operation further comprises maintaining one or more occurrence buffers to count occurrences of patterns.

14. (Original) The apparatus of claim 11, wherein the attribute association discovering operation further comprises pruning candidate patterns at a template level.

15. (Currently Amended) An article of manufacture for mining attribute associations in a relational data set, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

obtaining multiple items from the relational data set; and

discovering attribute associations using: (i) multi-attribute mining templates formed from at least a portion of the multiple items; and (ii) one or more mining preferences specified by a user;

wherein the multi-attribute mining templates are related by an anti-monotonicity property such that the property holds when mining top-down from k -itemsets to $(k + 1)$ -itemsets and when mining items defined by a set of k attributes to items defined by $k + 1$ attributes.

16. (Canceled).

17. (Original) The article of claim 15, wherein the one or more mining preferences specified by the user comprise specification of at least one of: (i) one or more desired multi-attribute mining templates; (ii) one or more irrelevant multi-attribute mining templates; and (iii) one or more rules concerning values of attributes in the multi-attribute mining templates.

18. (Original) The article of claim 15, wherein the attribute association discovering step further comprises generating candidate patterns at a template level.

19. (Original) The article of claim 18, wherein the candidate pattern generating step further comprises deriving candidate patterns of multi-attribute mining templates by merge-joining patterns of nodes of at least a portion of the templates without pre-sorting.

20. (Original) The article of claim 18, wherein the candidate pattern generating step further comprises maintaining one or more occurrence buffers to count occurrences of patterns.

21. (Original) The article of claim 18, wherein the attribute association discovering step further comprises pruning candidate patterns at a template level.